

1. (previously presented) A method of performing an application layer semantic analysis to detect information access anomalies, comprising:
 - a) capturing data packets;
 - b) filtering the captured data packets to detect information content;
 - c) processing packets based on semantics of an application or protocol;
 - d) generating a quantitative representation;
 - e) deriving a content signature from the quantitative representation;
 - f) deriving a prototypical model that includes a frequency view of a set of content signatures accessed by a given user, where the set of content signatures are indicative of content that is changing over time; and
 - g) detecting an application layer information access anomaly by using a semantic analysis to detect a given deviation from the prototypical model.
2. (previously presented) The method, according to claim 1, where the prototypical model also includes a time distribution of a set of content accesses by the given user.
3. (previously presented) The method, according to claim 1, where the prototypical model also includes a location distribution of a set of content accesses by the given user.
4. (previously presented) The method, according to claim 1, where the quantitative representation is captured as a content distribution vector that captures a frequency based distribution of key words in the message.
5. (previously presented) The method, according to claim 1, where the content signature is computed based on a moment statistic.
6. (previously presented) The method, according to claim 1, where the content

signature is computed as a hash of the information content.

7. (previously presented) The method, according to claim 1, where the content signature is computed via a document clustering technique where documents that share content signatures are classified together.

8. (previously presented) The method, according to claim 1, further including storing the information content, the content signature, and one or more attributes, where the attributes include one of: user identity, location of access, time of access, content type, content length, content hash, content encoding, and one or more content properties.

9. (previously presented) The method, according to claim 1, where mining is based on statistical clustering and distance based metrics.

10. (previously presented) The method, according to claim 9, where a statistical metric includes frequency of all content signatures accessed by a user.

11. (previously presented) The method, according to claim 9, where a statistical metric includes time of all content signatures accessed by a user.

12. (previously presented) The method, according to claim 9, where a statistical metric includes location of all content signatures accessed by a user.

13. (previously presented) The method, according to claim 1, where the prototypical model is derived by mining a content database.

14. (previously presented) The method, according to claim 1, where mining may be augmented by content aging, where information is periodically deleted from the content

database.

15. (previously presented) The method, according to claim 14, where content aging is a function of a mining algorithm and a type of information being monitored.

16. (previously presented) The method, according to claim 1, where the information access anomaly is based on a given user accessing given content from a given location at a given time.

17. (previously presented) The method, according to claim 16, where the information access anomaly is detected by a memory-based deviation where the given content accessed by the given user shows a deviation over normal content accessed.

18. (previously presented) The method, according to claim 16, where the information access anomaly is detected by a rare content condition, where the given user accesses given content that is rarely accessed by the given user.

19. (previously presented) The method, according to claim 16, where the information access anomaly is detected by a time deviation where the given user accesses the given content at a time different from historical access by the given user.

20. (previously presented) The method, according to claim 16, where the information access anomaly is detected by a location deviation where the given user accesses the given content from a location different from historical access by the given user.

21. (previously presented) The method, according to claim 1, further including processing the information access anomaly.

22. (previously presented) The method, according to claim 21, where processing the information access anomaly processing includes one of: positive correlation with at least one past security violation event, and negative correlation with a past false alarm or non-event.

23. (previously presented) The method, according to claim 1, where a set of consistent anomalies are classified into a pattern of misuse.

24. (previously presented) The method, according to claim 1, where the information access anomaly is detected in real-time.

25. (previously presented) The method, according to claim 1, where information access anomaly detection is used for real-time protection of information.

26. (previously presented) The method, according to claim 25, where real-time anomaly detection is used for protection via real-time alerts.

27. (previously presented) The method, according to claim 25, where real-time anomaly detection is used for real-time protection via denial of access.

28. (previously presented) The method, according to claim 25, where real-time anomaly detection is used for real-time protection via additional user validation.

29. (previously presented) The method as described in claim 1, where the data packets are associated with access to a confidential information repository.

30-41. (cancelled)

42. (previously presented) The apparatus of claim 49, implemented on a computing device and connected on a network as a passive tap.

43. (previously presented) The apparatus of claim 49, implemented as a network appliance that derives information transparently.

44. (previously presented) The apparatus of claim 49, implemented on an end-user computing device.

45. (previously presented) The apparatus of claim 49, implemented as a shim on an application server.

46. (cancelled)

47. (previously presented) The apparatus of claim 49, connected to an access control system to enable real-time monitoring of anomalous information access.

48. (previously presented) The apparatus of claim 49, configured to implement one or more compliance policies.

49. (previously presented) Apparatus, comprising:
a processor; and
a computer memory storing program instructions that when executed by the processor perform a method of detecting an information access anomaly, the method comprising:
monitoring data packets indicative of changing content over time;
generating a prototypical model; and
performing a semantic analysis against the prototypical model to identify an application level information access anomaly.
50. (cancelled)
51. (cancelled)
52. (new) The method of claim 1 wherein the application layer semantic analysis examines an entirety of an application layer without any application-specific limits.
53. (new) The apparatus as described in claim 49 wherein the detecting method examines an entirety of an application layer without any application-specific limits.